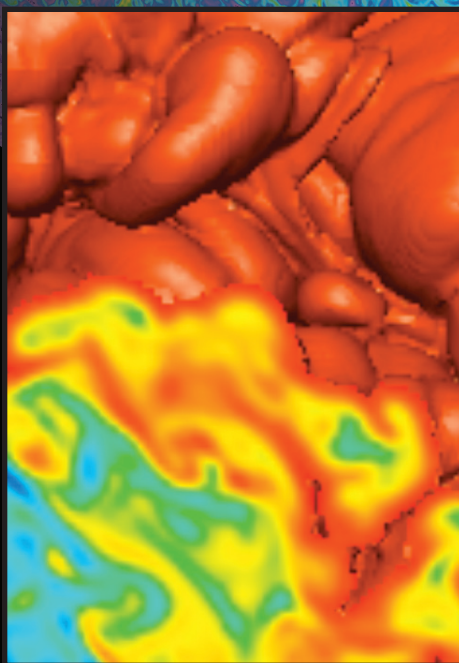


From Seeing to Understanding

Livermore computer scientists are revolutionizing the ways researchers visualize enormous amounts of supercomputer data.



Supercomputer simulations contain extraordinary amounts of detail. Livermore scientists are developing new methods for scientists to search images and locate the areas of interest. This image was generated on the Production Visualization Cluster—a visualization engine—using the MIRANDA code and Terascale Browser software. It depicts the instability of two fluids mixing together.

EXTRAORDINARILY complex, three-dimensional (3D) supercomputer simulations play a major role in making sure the nation's nuclear stockpile remains safe and reliable. The simulations require supercomputers performing trillions of operations per second (teraops) often for weeks at a time.

Understanding these simulations depends, to a great extent, on the human eye to carefully scrutinize the vast amounts of simulation information translated into still and moving images. These images allow researchers to gain insight into how a nuclear weapon operates and the effects of aging on its many components. Livermore computer scientists are developing new ways to see and understand the latest simulations by combining inexpensive and ubiquitous microprocessors, graphics cards favored by video-game fans, and open-source software. These components are the heart of the powerful visualization engines that turn reams of data into practical 3D images and movies.

Livermore's supercomputers dedicated to stockpile stewardship are part of the Advanced Simulation and Computing (ASC) Program. An element of the National Nuclear Security Administration (NNSA), ASC is advancing supercomputing so scientists can make the much higher fidelity physics and engineering simulations needed to assess the safety, reliability, and performance of the nation's nuclear weapons stockpile.

The 12.3-teraops White machine, in operation since 2000, is Livermore's most powerful ASC supercomputer. Two new ASC computers, Purple and BlueGene/L, will be delivered in 2005 and housed in Livermore's new \$91-million Terascale Simulation Facility (TSF). Purple will fulfill the ASC Program's long-sought goal of developing a machine that operates at 100 teraops, considered the entry point for prototype high-fidelity, full weapons system simulations. BlueGene/L, a research and evaluation machine, will have a peak performance of 180 to 360 teraops.

Modern ASC supercomputers, such as White, Purple, and BlueGene/L, consist of thousands of nodes, each composed of 2 to 16 microprocessors. These machines perform what is known as massively parallel computing in which nodes work together on a problem. They are also scalable; that is, simulations can be done with a few nodes or the entire set, and nodes can be added to tackle more difficult problems. (See *S&TR*, June 2004, pp. 12–20.)

The latest generation of ASC machines generate enormous amounts of data that are sometimes the result of weeks of round-the-clock number crunching. Three-dimensional, time-varying data sets of tens of terabytes (trillions of bytes) are now common, and petabyte (about 1,000 terabytes) data sets are on the horizon. As a result, an urgent need exists to develop new ways to visualize vast quantities of numbers.

Transforming Visualization

Livermore's Visual Interactive Environment for Weapons Simulation (VIEWS) Program, part of the Laboratory's ASC effort, is helping scientists visually explore, manage, and analyze data from advanced simulations. (See *S&TR*, October 2000, pp. 4–12.) The program is fulfilling a plan developed several years ago to transform the way scientists look at their data. This transformation is being accomplished by adopting new visualization computer architecture and developing software and analytical tools to run on the new hardware.

The VIEWS team's mantra is "see and understand." "Early on, when we were trying to display the results of our huge computer runs, the emphasis was on seeing a large amount of data at once," says computer scientist and VIEWS program leader Steve Louis. "We have progressed to displaying and managing the details of that data in high-resolution format."

Louis says the emphasis has shifted to *understanding* because "*seeing* only takes you so far. We want to make it easier to

find interesting regions in simulations and track those regions over time. We also want to visually compare a set of simulations or contrast data from a simulation with data from an experiment." The underlying goal is to support the ASC Program's vision of improved predictive capability for the performance of stockpile nuclear weapons and their components through experimentally validated simulation tools.

Currently, ASC supercomputers use visualization engines that turn the data produced by supercomputers into images displayed on individual computer monitors, larger-scale screens, or massive powerwalls. (See the **box** on p. 14.) These visualization engines and their systems software have until recently been supplied and integrated by commercial vendors. This approach worked well in the past, but it offers limited expansion capability because of the constraints of a shared-memory architecture. The processors and graphics cards used in the shared-memory architecture are proprietary and expensive.

Computer scientists use the term scalability to indicate the ability of a computer to handle larger and larger problems and data sets. "You can only scale so far with our present ASC visualization engines before the necessary hardware and time to run the simulations start getting very expensive," says Sean Ahern, a visualization project leader for VIEWS.

With an eye on the growing size of computer simulations, Livermore managers decided to transition from proprietary shared-memory visualization engines to groups or "clusters" of commercial personal computer (PC) microprocessors and high-end graphics cards found in gaming boxes and many PCs. When combined with a high-speed network running on a Linux operating system and software tools written in open-source (not vendor-proprietary) code, the clusters outperform the larger and significantly more expensive proprietary engines. The clusters are also easily expandable by simply adding more units.

A Closer Look at Clusters

Linux visualization clusters operate like modern supercomputers: They farm out problems in pieces to hundreds or thousands of microprocessors networked together and working in parallel. Clusters offer substantially more power in the same space—and at much less cost—than

the proprietary engines they replace. Individual cluster nodes typically have two microprocessors and one graphics card. These nodes have their own memories, in contrast to shared-memory designs.

The key to the visualization clusters' remarkable price–performance ratio is their high-end graphics cards containing

graphical processing units (GPUs). “The GPUs give us 10 times the performance for one-fifth the cost of cards found in previous ASC visualization engines,” says Ahern. Specialized, 3D GPUs were once available only in expensive workstations. Linux clusters now use gaming GPUs that cost between \$300 and \$400, are

Getting the Big Picture

To see the results of their simulations, Livermore researchers use a variety of display devices ranging from relatively small desktop monitors to powerwalls. Powerwalls work by aggregating, or “tiling,” the separate images from many projectors to create one seamless image. Large powerwalls exist in several Livermore buildings.

Powerwalls, which are typically the size of a conference room wall, allow a group of scientists to study still images or watch a movie, frame by frame. “Researchers can freeze images, pan, zoom,

move back and forth in time, and see details too subtle or small to discern on a desktop monitor,” says electronics engineer Bob Howe, head of infrastructure and facilities for the Visual Interactive Environment for Weapons Simulation (VIEWS) Program. At the same time, because of the powerwall's sheer size, users can still view the global problem while keeping the details in perspective. Powerwall displays are especially useful for presentations and formal reviews.

Livermore's new Terascale Simulation Facility (TSF) has two large powerwalls for major reviews and division meetings, one of which is used to display classified simulations in a room with removable seating for 130 people. A similar powerwall in a room with auditorium-style seating will be used for unclassified work. The TSF also has smaller powerwalls designed for more informal interactions.

Over the years, new products have been introduced to improve the resolution, clarity, and uniform brightness and color of powerwalls. Flexible screens have been replaced with hard, flat screens, and new projectors using digital-light-processing technology achieve higher contrast, greater brightness, and automated color balancing. For video delivery to powerwalls and other high-resolution displays, Howe is overseeing the transition from existing analog-based switching and delivery to newer digital technologies.

Some Livermore physicists have asked for stereoscopic capabilities to improve the three-dimensional (3D) information in powerwall presentations. Currently, 3D visualization is achieved by interactively shading and rotating an image to reveal the sides and details of objects slightly hidden behind foreground surfaces. Visualization experts plan to deploy active stereo technology, which uses high-frame-rate stereo projectors and requires viewers to wear shuttered goggles that repolarize about 45 frames per second per eye to minimize flicker.

Howe notes that although powerwalls have proven indispensable for presentations, scientists spend most of their time working in their offices alone or with a few colleagues. “Those scientists want larger displays and more pixels on their office machines, and we're working hard to provide that.” Howe and other VIEWS visualization experts are keeping a close eye on new monitor and projector designs, many of which are beginning to enter the consumer market.



A powerwall in Livermore's new Terascale Simulation Facility shows a simulation of results from an experiment mixing two liquids of different densities, which was conducted at the University of Arizona. Powerwalls work by aggregating, or “tiling,” the separate images from many projectors (inset) to create one seamless image.



powerful computers in their own right, and can calculate 100 billion operations per second (gigaops). Ever more powerful GPUs are announced every few months, and industry experts predict that video-game machines will contain GPUs capable of 1 teraops by 2006.

“Several years ago, we began watching the graphics cards appearing in PCs and gaming boxes. They didn’t have the performance we need, but we could see where the cards were heading,” says Louis. He notes that Livermore computer scientists, who were used to working closely with manufacturers in developing new computers and components, are now largely spectators in the multibillion-dollar gaming-hardware industry. Nonetheless, they have no objections to taking advantage of the hardware advances.

Paving the Way

The first Linux visualization cluster deployed at Livermore was the Production Visualization Cluster (PVC). PVC was designed to support unclassified applications on the 11.2-teraops Multiprogrammatic Capability Resource (MCR) machine and is being expanded to support the 22.9-teraops Thunder cluster supercomputer. With 64 nodes, each consisting of two processors and a graphics card, PVC went online in 2002.

By all measures, PVC has been highly successful. It is handling data sets of 23 terabytes to create animations involving 1 billion atoms. PVC generates these animations in about one-tenth the time and at one-fifth the cost of proprietary visualization engines, while simultaneously driving high-resolution displays in conference rooms and on powerwalls.

“PVC is our model for classified ASC visualization engines,” says Louis. The VIEWS team is preparing to deploy gViz, a 64-node cluster designed to support White, with each node consisting of two processors running at 3 gigahertz and sharing one graphics card. Similar clusters are planned to support Purple and BlueGene/L.

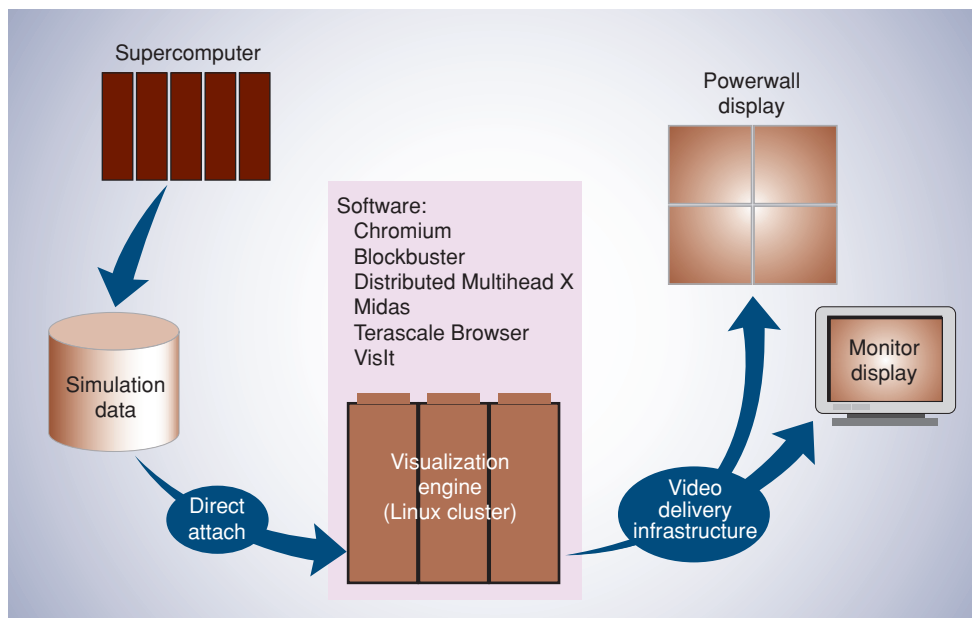
One important advantage of Linux cluster visualization engines is that clusters can be expanded easily. PVC is being tripled in size to support the unclassified demands brought on by Thunder. Similarly, gViz2 is a planned expansion of gViz to either 128 or 256 nodes.

Enormous Software Effort

The VIEWS team has overseen a huge shift in software, which has occurred simultaneously with the development and deployment of new visualization clusters. Many Livermore-developed user applications that ran on the old visualization systems were modified to run on the new clusters. Livermore software developers have also written new software, often with academic and industrial partners, to replace the proprietary software that shipped with the old machines. In one instance, the VIEWS team helped a graphics card manufacturer develop a software driver that allows the company’s GPUs to work together in the Linux operating system environment.



The visualization engine gViz is a 64-node Linux cluster designed to support Livermore’s White classified supercomputer. Each node contains two microprocessors that share one graphics card. Similar clusters are planned to support the Purple and BlueGene/L machines, which are scheduled to arrive in 2005.



Use of a visualization cluster involves interdependent software and hardware resources, including computational nodes, graphics services, display devices, and video switches.

The new cluster software is open source, which means that the source code—the software’s programming code—is freely available on the Internet through such organizations as SourceForge. “A benefit of this approach is that the public can use our software, make improvements, and notify us if they find any bugs,” says Ahern. Although several new software components are being developed under separate projects, many of the developers serve on multiple projects, thereby ensuring that all components work well with each other.

The software component Chromium provides a way for interactive 2D and 3D graphics applications to operate on clusters and allows the applications’ graphics cards to work together on a single visualization problem. Ahern and former Livermore computer scientist Randall Frank, in close collaboration with researchers from Stanford University, the University of Virginia, and Tungsten Graphics, designed the software. Chromium, which won an R&D 100

Award in 2004, supports any program that uses the OpenGL programming language, an industry-standard for drawing graphics. (See *S&TR*, October 2004, pp. 6–7.)

Ahern says, “Chromium is a Swiss army knife of graphics ‘tool kits’ because it fully exploits a cluster’s visualization capabilities.”

The Distributed Multihead X (DMX) software component combines multiple displays from multiple machines to create a single unified screen. It can create a display from two desktop machines or unify, for example, a 4-by-4 grid of displays (each attached to one of 16 computers) into one giant display. The software is particularly useful for powerwalls. DMX is bundled by Red Hat with its Linux software and is available as open-source software.

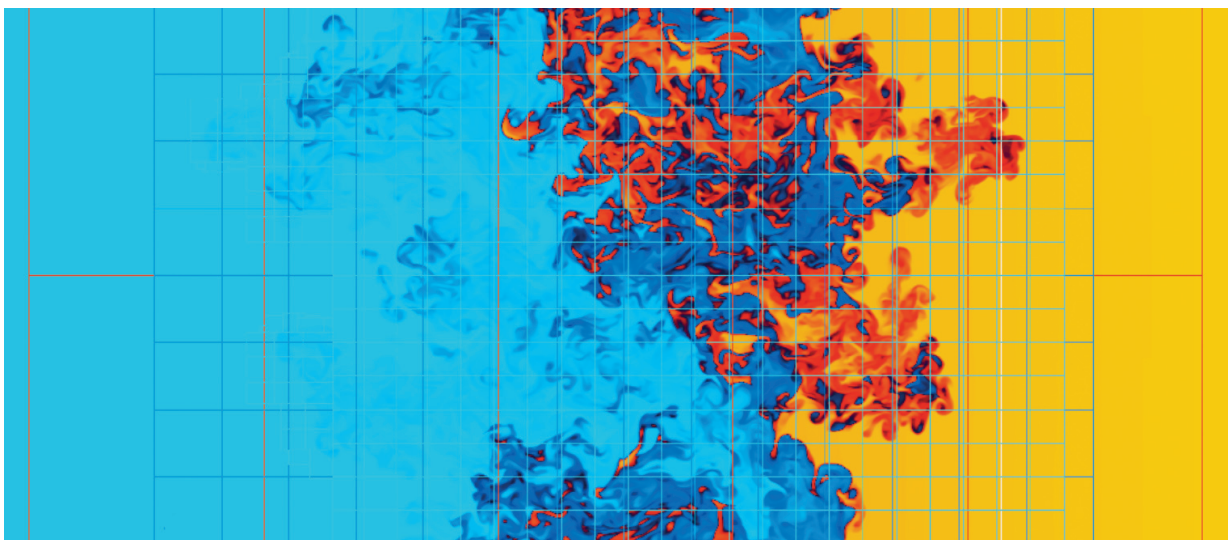
Physicists in Livermore’s Defense and Nuclear Technologies (DNT) Directorate use Livermore’s VisIt software tool to view extremely large scientific data sets. They can animate and manipulate the

visualizations and save them as images for presentations. VisIt has been modified to run on the new Linux clusters, and it is freely available.

Livermore’s Terascale Browser, a complementary program to VisIt, is another interactive tool that handles extremely large data sets and generates movies from the data. VIEWS developers are working on new ways for users to look for interesting areas with a coarser resolution and then zoom in with higher resolution to view the areas in finer detail.

Another Livermore program, MIDAS, permits visualizations of large simulations to run on office desktop monitors. MIDAS compresses images to one-tenth their original size without sacrificing the detail on the user’s display.

Two tools, Telepath and Blockbuster, streamline the process of using large displays. Telepath orchestrates a visualization session by automating the configuration of all the resources required for showing advanced simulations.



Researchers used the RAPTOR code to generate the data used by the software tool VisIt to create this two-dimensional image on the Production Visualization Cluster. The image shows the density differences between two gases mixing, which is caused by multiple shock-wave accelerations. The dense gas (sulfur hexafluoride) is depicted in gold, and the light gas (air) is depicted in light blue. The intermediate colors (darker blue to red) denote the mixing of the two gases on a small spatial scale. Overlaid on the figure is a grid that represents the data at increasingly finer spatial and temporal resolutions.

“Telepath greatly simplifies things for users,” says Ahern. Blockbuster, developed for clusters, plays high-resolution movies at 30 frames per second on powerwalls and other large displays.

Combing through a Billion Zones

As simulations grow in size, the ability of users to visually inspect their data has become increasingly limited. “We want to reduce the amount of information users must interact with so they can focus on the most useful details,” says Louis. “It would take months to look at all the data. We need ways to drill down and find the most important subsets.”

Livermore physicist Steve Langer studies how a laser beam interacts with a 2-millimeter-diameter stream of plasma. His simulations use a mesh composed of 12.7 billion zones, each of which depicts a different region of space. “That’s an enormous amount of information,” says Langer. “We can’t manually inspect 12.7 billion zones to find the ones in which interesting events are occurring.”

Given the sheer volume of data, computer science researchers need to help tease out the most relevant features. For example, Sapphire, a Livermore project, attempts to extract underlying patterns in a simulation. This research and development effort taps the field of data mining, which is the process of extracting useful information from raw data. The effort, led by Chandrika Kamath and funded by the Laboratory Directed Research and Development Program, uses such techniques as image processing and pattern recognition. Sapphire techniques are being applied to DNT data both to explore simulation results and to compare these results with experiments. (See *S&TR*, September 2000, pp. 20–22.)

The SimTracker data management tool helps scientists cope with organizing large amounts of simulation data using automated summaries. This tool, which originated at Livermore, has also been adopted by other national laboratories to archive, annotate, and share data. SimTracker summaries

allow users to easily access data analysis tools while browsing graphical snapshots, input and output files, and associated information, all tied into one convenient Web-based collection.

In addition to SimTracker, tools have been developed to automate manual data management processes and simplify the user interface to data. When combined with the suite of hardware and software visualization tools, these data management tools provide ASC users with what they need to manage, analyze, and present their data.

Verdict: Fast, Very Fast

The nearly unanimous opinion about the new Linux clusters is strong approval,

if not downright devotion. “Users are impressed with the clusters,” says computer scientist Hank Childs, who helps DNT physicists visualize complicated simulations on PVC for unclassified, stockpile stewardship–related work. “It’s a night-and-day difference between the clusters and the older shared-memory visualization engines. Visualization programs run 10 times faster.”

Ahern notes that the increased computational horsepower of the Linux clusters allows users to run larger simulations in the same amount of time and display simulations with greater resolution. Langer notes that the clusters are proving themselves especially adept at rotating images faster than the old machines.

Sapphire: Query Results

File Options

filename	Class	Distance	ART2D_12_3_0	ART2D_12_3_1	ART2D_1
	unknown	0.617259	0.426535	0.416434	0.0558684
	unknown	0.629357	0.426327	0.404481	0.0721123
	unknown	0.943996	0.380397	0.47637	0.0563651
	unknown	0.987242	0.418287	0.350253	0.0667956
	unknown	1.00731	0.553952		

Result 1 of 25.

Views

Show All Show Selected Plot

Search

New Refine

Sapphire: Similarity Based Object Retrieval

File

Filename /home/kamath/bookchapter/image1.fits Open

Use auto-generated features - variable tile size

Set Search Inputs v1.1

The Sapphire Similarity-Based Object Retrieval System allows a user to identify an object of interest in the simulation output (red box at right) and then return similar objects in the simulation database, ranked by degree of similarity (above).

Collaborations Key to Visualization Advances

In many respects, partnerships are key to the continuing success of the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing (ASC) Program. For example, NNSA managers and IBM and Livermore computer scientists have collaborated on the design of increasingly more powerful supercomputers.

The Academic Strategic Alliances Program, an ASC program, engages the best minds in the U.S. academic community to help advance simulation science. That goal is shared by Livermore's Institute for Scientific Computing Research. Each year, the institute brings visiting postdoctoral researchers, faculty, and graduate students to the Laboratory. It also hosts an increasing number of undergraduate students majoring in computer science, who participate in computer programming internships.

Most visitors are integrated into the Center for Applied Scientific Computing (CASC), the research arm of Livermore's Computation Directorate, where they work on high-profile research projects. CASC has long-term visualization research relationships with the University of California (UC) at Davis, Duke University, Georgia Institute of Technology (Georgia Tech), University of North Carolina at Chapel Hill, University of Utah, and Stanford University.

Many people first become acquainted with Livermore visualization research at Laboratory-sponsored workshops. "We sponsor workshops so that a number of knowledgeable people can gather to think through issues, such as what it takes to support a 100-teraops machine," says computer scientist Mark Duchaineau, who oversees many students working in CASC.

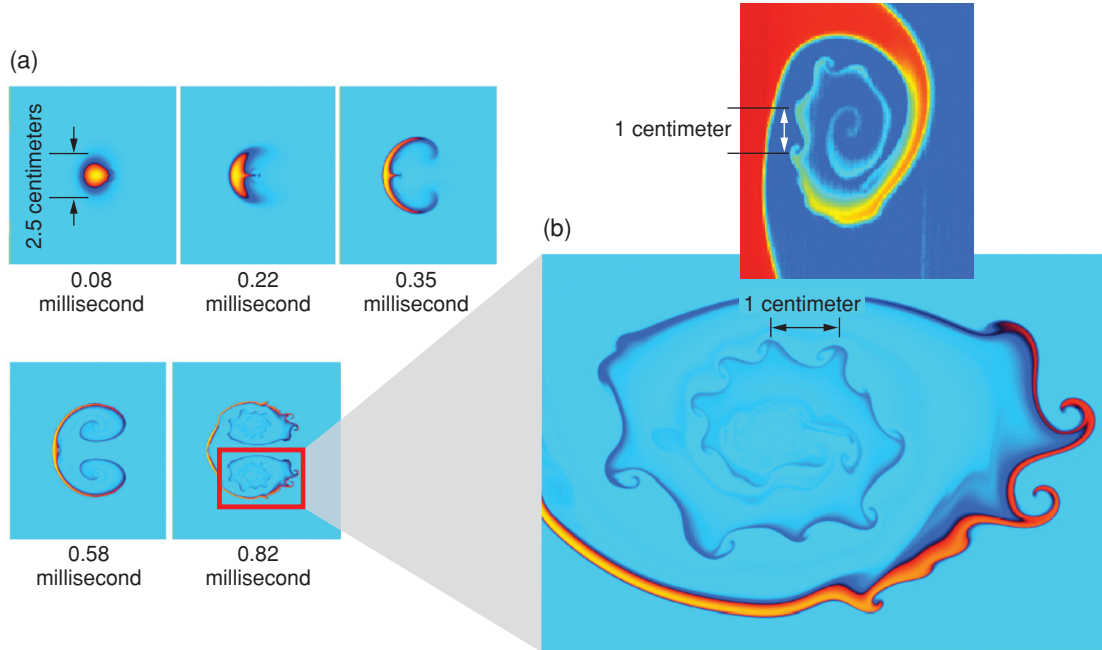
Duchaineau notes that visiting faculty and students gain valuable experience by working with some of the world's most powerful supercomputers. "There's an atmosphere here that is hard to duplicate at other institutions," says Duchaineau. "Visiting students and faculty become part of the big science we do here. They have access to some of the best equipment in the world, and they assist us with visualizing such phenomena as material properties, shock waves, hydrodynamics, radiation transport, and astrophysics. Our research team publishes extensively, and that's good for people's careers."

Computer scientist Daniel Laney obtained his Ph.D. from the UC Davis Department of Applied Sciences at Livermore. Laney's thesis, completed at CASC, focused on novel ways to compress data. He chose Livermore because of the variety of the Laboratory's physics projects and the presence of powerful supercomputers. He is currently working on new ways to model planned experiments on Livermore's National Ignition Facility. "In CASC, we provide the proof of concept and then depend on VIEWS software engineers to make it a practical application for the end user," he says.

Computer scientist Peter Lindstrom works with both students and faculty from Georgia Tech and the University of North Carolina at Chapel Hill, two of the nation's top universities in supercomputing visualization. Both schools also have remote access to some of Livermore's unclassified supercomputers. "The payoff for Livermore," says Lindstrom, "is access to some of the best people in the visualization community."

Computer scientist Mark Duchaineau works in front of a powerwall depicting a lattice of 1 billion copper atoms undergoing enormous strain. Duchaineau helps visiting students and faculty members make the most of their supercomputer visualization research at Livermore.





(a) These images depict what happens over time when a shock wave accelerates a quantity of sulfur hexafluoride contained in a cylinder diffused with ambient air. The shock wave causes the gas to spiral, and the spirals form tiny unstable vortices. The images were created using the RAPTOR code on the Production Visualization Cluster. (b) The larger image depicts a magnified view of a spiral. The result from an experiment conducted at the University of Arizona (inset) is similar to the simulation result in the larger image.

With plans well under way to bring gViz online and retire the old visualization engines, Louis and other VIEWS managers are looking ahead to purchasing visualization clusters to support Purple and Blue Gene/L. At the same time, computer scientists are searching for ways to make the clusters process data more efficiently. “We know we’re not yet taking full advantage of the Linux clusters, especially the graphics cards,” says Louis. GPUs are so powerful that the VIEWS team and others are exploring their potential for general-purpose computing.

The VIEWS Program (recently renamed Data and Visualization Sciences) is also seeking a hardware solution to compositing. The compositing process pieces together bits of an image, each done by a separate node, into a whole. Currently performed by software, the

technique could be made faster if done by a specialized card linked to each GPU.

Louis says the many advances taking place in hardware and software are permitting researchers to not only see their simulations in breathtaking detail but also understand them to a much greater degree. The winners are Livermore stockpile stewardship scientists and, ultimately, national security.

—Arnie Heller

Key Words: Advanced Simulation and Computing (ASC) Program, BlueGene/L, Center for Applied Scientific Computing (CASC), Chromium, Distributed Multihead X (DMX), gViz, Linux, Purple, stockpile stewardship, supercomputing, Terascale Browser, Visual Interactive Environment for Weapons Simulation (VIEWS), VisIt, White.

For further information contact Steve Louis (925) 422-1550 (louis3@llnl.gov).